

# Discovering Neuronal Cell Types and Their Gene Expression Profiles Using a Spatial Point Process Mixture Model

Furong Huang

UC Irvine

furongh@uci.edu

Animashree Anandkumar

UC Irvine

a.anandkumar@uci.edu

Christian Borgs

Microsoft Research

Christian.Borgs@microsoft.com

Jennifer Chayes

Microsoft Research

jchayes@microsoft.com

Ernest Fraenkel

MIT

fraenkel@mit.edu

Michael Hawrylycz

Allen Institute

MikeH@alleninstitute.org

Ed Lein

Allen Institute

EdL@alleninstitute.org

Alessandro Ingresso

Politecnico di Torino

alessandro.ingrosso@polito.it

Srinivas Turaga

HHMI Janelia Research Campus

turagas@janelia.hhmi.org

## Abstract

Cataloging the neuronal cell types that comprise circuitry of individual brain regions is a major goal of modern neuroscience and the BRAIN initiative. Single-cell RNA sequencing can now be used to measure the gene expression profiles of individual neurons and to categorize neurons based on their gene expression profiles. While the single-cell techniques are extremely powerful and hold great promise, they are currently still labor intensive, have a high cost per cell, and, most importantly, do not provide information on spatial distribution of cell types in specific regions of the brain. We propose a complementary approach that uses computational methods to infer the cell types and their gene expression profiles through analysis of brain-wide single-cell resolution in situ hybridization (ISH) imagery contained in the Allen Brain Atlas (ABA). We measure the spatial distribution of neurons labeled in the ISH image for each gene and model it as a spatial point process mixture, whose mixture weights are given by the cell types which express that gene. By fitting a point process mixture model jointly to the ISH images, we infer both the spatial point process distribution for each cell type and their gene expression profile. We validate our predictions of cell type-specific gene expression profiles using single cell RNA sequencing data, recently published for the mouse somatosensory cortex. Jointly with the gene expression profiles, cell features such as cell size, orientation, intensity and local density level are inferred per cell type.

## 1 Introduction

### 1.1 Motivations and Goals

The human brain comprises about one hundred billion neurons and one trillion supporting glial cells. These cells are specialized into a surprising diversity of cell types. The retina alone boasts well over 50 cell types, and it is an active area of research to perform a census of the various neuronal cell types that comprise the central nervous system. Many criteria have been used to categorize neuronal cell types, from neuronal morphology and connectivity to their functional response properties. Neurons can also be categorized based on the proteins they make. Immunohistochemistry has been used with great success for many decades to differentiate excitatory neurons from inhibitory neurons by labeling for known proteins involved in the synthesis and regulation of glutamate and GABA, the primary excitatory and inhibitory neurotransmitters respectively.

More recently, there has been an effort to systematically measure the complete transcriptome of single neurons. Single-cell RNA sequencing (RNA-Seq) is an extremely powerful technique that can quantitatively determine the expression level of every gene that is expressed in individual neurons. This so-called transcriptome or gene expression

/ transcription profile can then be used to define cell types by clustering. A recent study produced the most comprehensive census of cell types to date in the mouse somatosensory cortex and hippocampus by performing single-cell RNA-Seq on over 3000 neurons [12]. While this study is quite exciting, trying to replicate it for all brain regions might well require the equivalent of a thousand such experiments. Thus, it is likely that the unprecedented insights that RNA-Seq can provide will be slow to arrive. More importantly, single cell sequencing methods are not currently able to capture the precise three-dimensional location of the individual neurons.

Here we propose a complementary approach that uses computational strategies to identify cell types and their spatial distribution by re-analysing data published by the Allen Institute for Brain Research. The Allen Brain Atlas (ABA) contains cellular resolution brain-wide *in-situ* hybridization (ISH) images for 20,000 genes<sup>1</sup>. ISH is a histological technique that labels the mRNA in all cells expressing the corresponding gene in a manner roughly proportion to the gene expression level. An example of an ISH image can be seen in figure 1(a).

The ABA contains genome-wide and brain-wide ISH images of the adult mouse brain. These images were generated by slicing the brain into a series of 25  $\mu\text{m}$  thin sections and performing ISH. Image series of ISH performed for different genes come from different mouse brains, since ISH can only be performed for one gene at a time. The ISH image series for different genes were then computationally aligned into a common reference brain coordinate system. Such data have been productively used to infer the average transcriptomes corresponding to different brain regions.

It is commonly thought that the ABA cannot be used to infer the transcriptomes of individual cells in a given brain region since mouse brains cannot be aligned to the precision of a single cell. This is because there is individual variation in the precise number and location of neurons from brain to brain. However, we expect that the average number and spatial distribution of neurons from each cell type to be conserved from brain to brain, for a given brain area. More concretely, we might expect that parvalbumin-expressing (PV) inhibitory interneurons in layer 2/3 of the mouse somatosensory cortex comprise approximately 7% of all neurons and have a conserved spatial and size distribution from brain to brain. We use this fact to derive a method for simultaneously inferring the cell types in a given brain region and their gene expression profiles from the ABA.

We propose to model the spatial distribution of neurons in a brain as being generated by sampling from an unknown but consistent brain-region and cell-type dependent spatial point process distribution. And since each gene might only be expressed in a subset of cell types, an ISH image for a single gene can be thought of as a mixture of spatial point processes where the mixture weights represent the individual cell types expressing that gene. We infer cell types, their gene expression profiles and their spatial distribution by unmixing the spatial point processes corresponding to the ISH images for 1743 genes. This is in notable contrast to the information provided by single-cell RNA sequencing which can only measure the gene expression profile of individual cells to high accuracy but where, due to the destructive measurement process, all information about the spatial position and distribution of cell types is lost.

## 1.2 Previous Work

Allen Brain Atlas (ABA) [9] is a landmark study which mapped the gene expression of about 20,000 genes across the entire mouse brain. The ABA dataset consists of cellular high-resolution 2d imagery of *in-situ* hybridized series of brain sections, digitally aligned to a common reference atlas. However, since the *in-situ* images for each gene come from different mouse brains and since there is significant variability in the individual locations of labeled cells, it is not possible to register brain-wide gene expression at a resolution higher than about 250  $\mu\text{m}$ . Therefore, the cellular resolution detail was down-sampled to construct a coarser 3d representation of the average gene expression level in 250  $\mu\text{m} \times 250 \mu\text{m} \times 250 \mu\text{m}$  voxels.

The coarse-resolution averaged gene expression representation has been widely used and analyzed to understand differences in gene expression at the level of brain region. Hawrylycz et al [5] analyzed the correlational structure of gene expression at this scale, across the entire mouse brain. However, due to the poor resolution of the average gene expression representation, it has proven challenging to use the ABA to discover the microstructure of gene expression within a brain region. To address this issue from a complementary perspective, Grange et al [3] used the gene expression profiles of 64 known cell-types, combined with linear unmixing to determine the spatial distribution of

<sup>1</sup> Although the Atlas contains ISH data for approximately 20,000 distinct mouse genes, we focus on the top 1743 reliable genes whose sagittal and coronal experiments are highly correlated.

these known cell-types. However, such an approach can be confounded by the presence of cell-types whose expression profiles have yet to be characterized, and limited by the resolution of the averaged gene expression representation.

In contrast to previous approaches, we aim to solve the difficult problem of automatically discovering the gene expression profiles of cell-types within a brain region by analyzing the original cellular resolution ISH imagery. We propose to use the spatial distributions of labeled cells, and their shapes and sizes, which are a far richer representation than simply the average expression level in  $250\mu m \times 250\mu m \times 250\mu m$  voxels. This spatial point process is then un-mixed to determine the gene expression profile of cell types.

Most previous work on unmixing point process mixtures adopted parametric generative models where the point process is limited to some distribution family such as Poisson or Gaussian [6, 7]. However, since we are not interested in building a generative model of a point process, but rather care more about inferring the mixing proportions (gene expression profile), we take a simpler parameter-free approach. This approach models only the statistics of the point process, but is not a generative model, and so cannot be used to model individual points/cells.

## 2 Modeling the Spatial Distribution of Cell-types Using Spatial Point Process Features

Most analyses of the ABA *in situ* hybridization dataset have utilized a simple measure of average expression level in relatively large  $250\mu m \times 250\mu m \times 250\mu m$  voxels of brain tissue. Due to the large volume over which the expression level is averaged, such a representation cannot distinguish between large numbers of cells expressing small amounts of RNA vs. small numbers of cells expressing large amounts of RNA. All information about the spatial organization of labeled cells, their shapes, sizes and spatial density are lost and summarized by a single scalar number. Here, we describe a more sophisticated representation of the labeled cells in an ISH image based on marked spatial point processes.

### 2.1 The Marked Spatial Point Process Representation of ISH Images

Our approach requires processing the high-resolution ISH images to detect individual labeled cells and their visual characteristics. We developed a cell detection algorithm described in the Supplementary section. Our algorithm additionally also estimates the expression level of each detected cell, its shape, size and orientation. Figure 1(a) and Figure 1(b) illustrate the results of our cell detection algorithm.

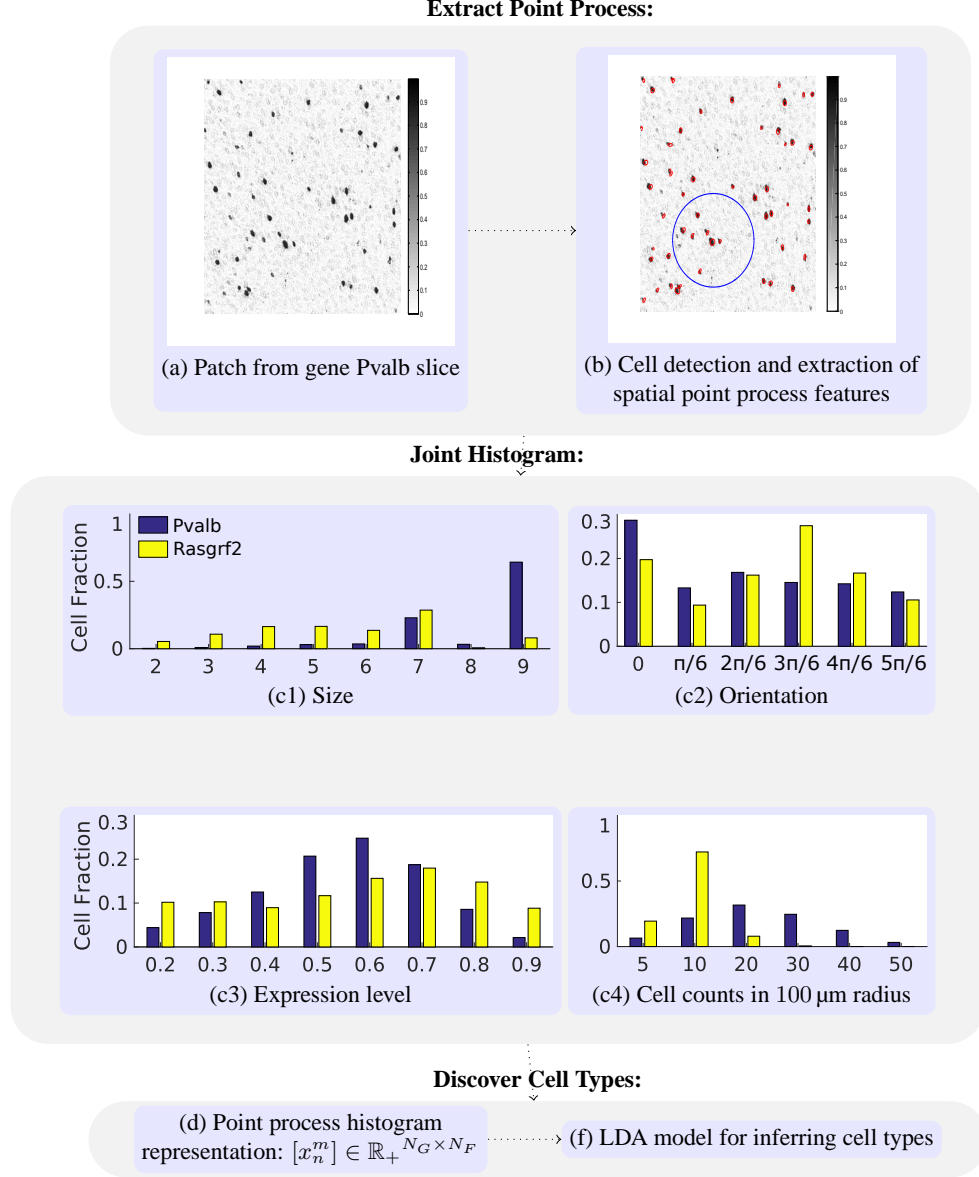
Since cell-types differ not only in terms of gene expression pattern, but also display a diversity of shapes, sizes and spatial densities, we sought to characterize these properties. We measured: (1) **cell size**  $s = [r_1, r_2]$ : the radius in two principal directions of an ellipse fit to each cell; (2) **cell orientation**  $\phi$ : the orientation of the first principle axis of the ellipse; (3) **gene intensity level**  $p$ : intensity of labeling of a cell relative to the image background; (4) **spatial distribution**  $c$ : the number of cells within a local area centered around the cell, which can be regarded as a measure of the local cell density.

The collection of detected cells within an atlas-defined brain region, along with their features, constitutes a marked spatial point process. This point process is considered “marked”, because each point is characterized by the shape, size, expression level and local density features, in addition to just their location in space.

### 2.2 A Model-free Approach to Representing Spatial Point Processes Using Joint Feature Histograms

The statistical modeling of repulsive spatial point processes such as those that arise in biology is non-trivial, and many generative models such as determinantal point processes [8] and Matern point processes have high computational complexity. But since we are not interested in directly modeling the individual labeled cells, but instead in modeling only their aggregate spatial statistics, and in inferring their gene expression profiles, we can take a simpler approach.

We use a *joint histogram* simple statistics of the collection of detected cells to characterize the underlying point process from which they are drawn. This is an empirical moment approach which side-steps the need to carefully define a generative point process distribution.



**Figure 1:** Overview of the proposed framework - Discovering Neuronal cell Types via Un-mixing of Spatial Point Process Mixtures. (a) & (b) An *in situ* hybridization image for gene Pvalb along with detected cells. (c) Marginalized point process feature histograms for genes Pvalb and Rasgrf2. Note that size denotes the principal axis diameter. We have  $N_G$  genes and 4d joint histogram with  $N_F$  bins.

As we describe in the next section, we propose to model the point process measured from the ISH image for each gene as a mixture of point processes belonging to individual cell-types. For this, we use a linear mixing model, the Latent Dirichlet Allocation model. The use of this model is greatly simplified if we carefully choose our feature representation such that the linear mixture of point processes results in a linear mixture of histogram statistics. This is clearly the case for the features we have chosen. For instance, if we sample equally from two point process distributions  $P_1$  and  $P_2$  with average densities of  $d_1$  and  $d_2$ , the addition of these two point processes  $P = P_1 + P_2$  results in the addition of the two densities  $d = d_1 + d_2$ . This is not the case for second order features, such as the distances to the nearest neighbors, which would have a more nonlinear relationship.

In figure 1(c), we display marginal histograms corresponding to the joint histogram for two genes, Pvalb and

Rasgrf2, which are well-known markers for a specific class of inhibitory and excitatory cortical neuronal cell-types respectively.

### 3 Un-mixing Spatial Point Processes to Discover Cell-types

#### 3.1 Generative Model: A Variation of Latent Dirichlet Allocation

The spatial point process histogram representation of the ABA ISH dataset results, for each brain region, is an  $N_F \times N_G$  matrix  $[x_n^m]$ , where  $N_F$  is the total number of histogram bins (henceforward called the number of histogram features)<sup>2</sup>,  $N_G$  is the number of genes, and  $x_n^m$  is the number of cells expressing gene  $n$  in histogram bin  $m$ .

We model the gene-spatial histogram matrix  $[x_n^m]$  by assuming it is generated by a Variation of Latent Dirichlet Allocation (vLDA) [2] model of cell types. This matrix factorization based latent variable model assumes that the ISH histograms are generated from a small number of cell-types,  $K$ , and each cell-type  $i$  is associated with a type-dependent spatial point process histogram  $h_i$  and a gene expression profile  $\beta_i$ .

Our generative model for each histogram bin  $m$  (characterizing a particular bin in the size/ orientation/ gene profile/ spatial distribution) is as follows: Let  $L^m = \sum_n^{N_G} x_n^m$  be the detected number of cells in the joint histogram bin  $m$ . For each cell  $l$  in this bin, its cell-type  $t$  is sampled from the multinomial distribution  $h^m$ . And given the cell-type  $t$  of cell  $l$ , the genes  $n$  expressed by this cell are sampled from a multinomial distribution given by the type-dependent gene expression profile/distribution  $\beta^t$ . For a given gene  $n$  and histogram bin  $m$ , this generative process determines the number of cells that would be detected  $x_n^m$ .

We further place a Dirichlet prior over  $h^m \sim \text{Dir}(\alpha)$ , with the concentration parameter  $\alpha$  which determines the prior probability over the number of cell-types present in a given histogram bin  $m$ . This prior represents our prior knowledge of how many cell-types express each gene, and also how well our feature representation separates cells of different types into different histogram bins. In principle, we could generalize this to be a gene-specific prior, if we had such information available. We could also use  $\alpha$  to incorporate information about our prior knowledge over the distribution of cells from each cell-type, for instance that excitatory neurons greatly outnumber inhibitory neurons in a roughly 5 : 1 ratio.

We now describe how we estimate the model parameters – the cell-type specific multinomial gene expression profile  $\beta$  and the cell-type specific spatial point process histogram  $h$  from the gene-specific spatial point process histograms measured from the ISH images.

#### 3.2 Estimating the Cell-type Dependent Gene Expression Profile $\beta$

After testing several estimation methods for the parameters of our model, we found that non-negative matrix factorization (NMF) performed well in estimating the cell-type specific gene expression profiles  $\beta$ , see Figure 2a. We solve the following optimization problem:

$$\min_{\beta, h} \sum_m^{N_F} \sum_n^{N_G} (x_n^m - \sum_t^K h_t^m \beta_n^t L^m)^2, \quad s.t. \quad \beta_n^t \geq 0, \sum_n^{N_G} \beta_n^t = 1, h_t^m \geq 0, \sum_t^K h_t^m = 1 \quad (1)$$

Here, the non-negativity and sum-to-one constraints on  $h_t^m$  and  $\beta_n^t$  ensure that  $h$  and  $\beta$  result in properly normalized multinomial distributions. While this estimation procedure results in joint estimates for  $h$  and  $\beta$ , it does not enforce the Dirichlet prior over  $h$ . So we refine our NMF-derived estimates for  $h$  using variational inference [2].

<sup>2</sup>Note that there are two types of *features* – the features characterizing each detected cell, and the features characterizing the collection of detected cells that constitute a single sample from a spatial point process

### 3.3 Estimating the Cell-type Dependent Spatial Point Process Histogram $h$

We use a standard maximum likelihood estimation procedure for  $h$  [2]. Iteratively, we refine the inference of the cell type membership  $h^m \in \Delta_K$  under each joint histogram feature  $m$ . We update  $h_i^m$  until convergence [11].

$$h_i^m \leftarrow \frac{1}{L^m + \sum_t^K \alpha_t} \sum_{n=1}^{N_G} x_n^m \frac{h_i^m \beta_n^i}{\sum_{l=1}^K h_l^m \beta_n^l} + \alpha_i, \forall i \in [K], m \in [N_F] \quad (2)$$

Recall that the Dirichlet prior  $\alpha$  encodes the number of cell-types that we expect on average to express each gene. We set  $\alpha$  to be a symmetric Dirichlet with  $\alpha_1 = \alpha_2 = \dots = \alpha_K$ , and  $\sum_t \alpha_t = 0.01$  for all cell-types  $t$ . In practice, we observe that our estimates of  $h$  are fairly insensitive to the specific choice for  $\alpha$  as long as  $\sum_t \alpha_t$  is small enough. The smaller  $\alpha$  is, the fewer cell-types expressing a given gene we expect to observe in a single histogram bin.

## 4 Results and Evaluation

### 4.1 Implementation Details

We tested our proposed cell-type discovery algorithm using the high-resolution *in situ* hybridization image series for 1743 of the most reliably imaged and annotated genes in the ABA. Individual cells were detected in the cellular resolution ISH images using custom algorithms (detailed in Supplementary Information). For each detected cell, we fit ellipses and extract several local features: (a) size and shape represented as the diameters along the principle axes of the ellipse, (b) orientation of the first principle axis, (c) gene intensity level as measured by the intensity of labeling of the cell body, and (d) the number of cells detected with-in a  $100 \mu m$  radius around the cell, which is a measure of the local cell density. We aligned the ISH images to the ABA reference atlas and, for this paper, focused our attention on cells in the somatosensory cortex, since independent RNA-Seq data exist for this region the can be used to evaluate our approach. We computed joint histograms for the collection of cells found with-in the somatosensory cortex, resulting in a spatial point process feature vector of  $N_F = 10010$  histogram bins per gene.

**Synthetic experiment:** The vLDA model we proposed is then fit to  $N_G \times N_F$  gene point process histogram matrix to estimate the cell-type gene expression profile matrix  $\beta$  using the non-negative matrix factorization (NNMF) algorithm. The reason why we choose NNMF over Variational Inference (which is a popular approach for LDA) for  $\beta$  estimation is that NNMF produces more accurate  $\beta$  estimation in simulated data, illustrated in Fig 2a. In the synthetic experiment, we simulate point process data ( with some predefined golden standard  $\beta$ ) and use the data to estimate  $\hat{\beta}$ . The errors were computed after pairing the estimated columns of  $\beta$  with a closest golden standard  $\beta$  column via hypothesis testing. Note that the columns of  $\beta$  are normalized to 1, so the errors are bounded.


### 4.2 Evaluating Cell-type Gene Expression Profile Predictions

A recent study performed single-cell RNA sequencing on 1691 neurons isolated from mouse somatosensory cortex. We use this dataset to evaluate the quality of the cell-types we discover.

The single cell RNA-seq data,  $G := [g^1 | g^2 | \dots | g^{N_C}] \in \mathbb{R}^{N_G \times N_C}$ , contains the gene expression profiles for  $N_C = 1691$  cells. We infer the cell types  $h^i$  for these cells using equation (2), and then compute the likelihood  $L^i$  of observing each for each cell under our estimated cell-type dependent gene expression profile matrix  $\beta$  using equation (4). We can then evaluate the perplexity, a commonly used measure of goodness of fit under the vLDA model, of single cell RNA-seq data on the model we learned from our spatial point process data.

The perplexity score is a standard metric, which is defined as the geometric mean per-cell likelihood. It is a monotonically decreasing function of the log-likelihood  $\mathcal{L}(G)$  of test data  $G$ .

$$\text{perplexity}(G) = \exp\left(-\frac{\sum_{i=1}^{N_C} \log p(g^i)}{\sum_{i=1}^{N_C} L^i}\right) \quad (3)$$



**Figure 2:** (a) Synthetic Experiment : comparison of Non-negative Matrix Factorization (NNMF) with Variational Inference (VI) on simulated point process cell data using known gene expression profile  $\beta$ . An additional robustness test of NNMF is done to see how good the algorithm is when a wrong number of cell types  $K$  is input. A permutation test (shuffling the gene expression levels between cell) is done to access statistical significance. Comparing with permute test shows that our cell-types are significantly different from chance. Error per type is computed by pairing the columns of estimated  $\hat{\beta}$  with the columns of the ground-truth  $\beta$ . (b) Comparison of gene expression profiles recovered for cell-types in the somatosensory cortex by fitting an LDA model using spatial point process features (ours) vs the standard average gene expression level feature (baseline). Our features provide a significantly better match, with lower perplexity, to ground truth single-cell RNA sequencing derived transcriptomes. A permutation test is done to access statistical significance. Perplexity is computed by matching to surrogate single-cell RNA transcriptomes by shuffling the gene expression levels between cells. Comparing with permute test shows that our cell-types are significantly different from chance.

where the likelihood is evaluated as


$$p(g^m | h^m, \alpha, \beta) = \frac{\Gamma(\sum_i \alpha_i)}{\prod_i \Gamma(\alpha_i)} \prod_{i=1}^k (h_i^m)^{\alpha_i - 1} \prod_{j=1}^{L^m} \left( \sum_{i=1}^k \sum_{n=1}^{N_G} \delta_{g_j^m, e^n} h_i^m \beta_n^i \right). \quad (4)$$

where  $\delta_{i,j}$  is the Kronecker delta,  $\delta_{i,j} = 1$  when  $i = j$  and 0 otherwise.  $e^n$  is the  $n^{\text{th}}$  basis vector.

### 4.3 Comparison to Standard Average Gene Expression Features Baseline and a Permutation Test for Significance

Here we demonstrate the superiority of our method and its statistical significance in two ways. First we compared the perplexity of the single-cell RNA seq dataset  $G$  under our model (figure 2b, solid blue) against the perplexity of a surrogate dataset with the same marginal statistics, but whose gene-cell correlations were destroyed (figure 2b, dashed blue). We generated this surrogate dataset by randomly permuting the gene expression levels for each gene across cells. This permuted dataset had a significantly higher (worse) perplexity than the true single-cell dataset. This demonstrates that our model trained to un-mix the ISH-derived spatial point processes discovered cell-types whose gene expression profiles are significantly better match to single-cells than by chance.

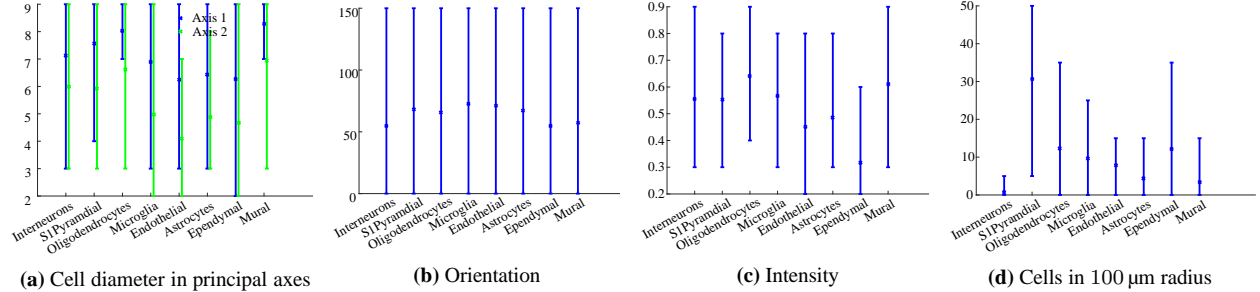
We also compared the predictions of cell-type gene expression profiles derived by un-mixing our spatial point process features against gene expression profiles derived by un-mixing the more standard  $250\mu\text{m} \times 250\mu\text{m} \times 250\mu\text{m}$  averaged gene expression level features. We see a very large improvement in perplexity by switching from the standard simple averaging of gene expression, to extracting spatial point process features (figure 2b). The single-cell RNA seq dataset analysis from figure 2b shows that the perplexity of our recovered cell-types rapidly flattens after we recover approximately 10 clusters ( $K = 10$ ).



**Figure 3:** Estimated memberships  $\beta$  on marker genes for 8 cell types. These marker genes are used to label the columns of the membership matrix.

#### 4.4 A Brief Analysis of Recovered Cell Types in Somatosensory Cortex

In this section we describe the representative spatial point process statistics and gene expressions for 8 cell-types we recovered. We attempted to align our 8 clusters to cell-types defined by [12] in the single-cell RNA sequencing paper. We found high overlap in the gene expression profiles for all 8 clusters with known cell-types defined in [12], *Interneurons*, *S1 Pyramidal*, *Mural*, *Endothelial*, *Microglia*, *Ependymal*, *Astrocytes* and *Oligodendrocytes*, in Figure 3.



**Figure 4:** Figure of 5% and 95% percentile estimated cell features for 8 cell types we detected. Inference is performed on the Spatial point process histograms data we estimated.

The estimate of  $\beta$  was combined with MLE to infer the cell-type specific spatial point process representation  $h_l^m$ . In examining the spatial point process distributions that we predict for each of these cell types, we discover that while the distribution of cell body orientations is quite broad and similar across cell types, the cell count distribution, which is a measure of cell density, varies in a systematic way from one cell type to another. Fig 4d shows that inhibitory Interneurons are less dense than S1Pyramidal neurons. This is consistent with their known prevalence, roughly 20% of all neurons are GABAergic interneurons [10], while the remaining 80% are excitatory glutamatergic pyramidal neurons. As expected, this excitatory neuronal category of S1Pyramidal is the most common and hence most dense class of neuronal cells. They also have slightly larger cell bodies, compared to interneurons, as can be seen in Fig 4a. The remaining 6 cell types correspond to various glial sub-types.



## 5 Conclusion

We developed a computational method for discovering cell types in a brain region by analyzing the high-resolution *in situ* hybridization image series from the Allen Brain Atlas. Under the assumption that cell types have unique spatial distributions and gene expression profiles, we used a varied latent Dirichlet allocation (vLDA) based on spatial point process mixture model to simultaneously infer the cell feature spatial distribution and gene expression profiles of cell types. By comparing our gene expression profile predictions to a single-cell RNA sequencing dataset, we demonstrated that our model improves significantly on state of the art.

The accuracy of our method relies heavily on the assumption that cell-types differ in their spatial distribution, and that our point process features perform a good job of distinguishing these differences. Thus the performance of our method can be improved by better estimates of better features. We would expect our method to perform better for large brain areas, which can be more accurately aligned, and which have more cells to estimate point process features.

There are several modifications to our vLDA model which might improve the faithfulness of our generative model to the biology. We place a symmetric Dirichlet prior over cell-type multinomial distribution  $h^m$  for a given histogram bin  $m$ . This assumes that the number of cell-types expressing each gene is the same for all genes. But since some genes are expressed more commonly and non-specifically than others, we might expect a gene-specific prior to be a better model. Further, the symmetric Dirichlet assumes that all cell-types have equal proportions of cells. But evidence suggests that excitatory neurons are more common than inhibitory neurons in cortex [4], and using a non-uniform Dirichlet prior could account for this.

## References

- [1] Website: 2014 allen institute for brain science. allen mouse brain atlas [internet]. Available from: <http://mouse.brain-map.org/>. Accessed: 2014-11-06.
- [2] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003.
- [3] Pascal Grange, Jason W Bohland, Benjamin W Okaty, Ken Sugino, Hemant Bokil, Sacha B Nelson, Lydia Ng, Michael Hawrylycz, and Partha P Mitra. Cell-type-based model explaining coexpression patterns of genes in the brain. *Proceedings of the National Academy of Sciences*, 111(14):5397–5402, 2014.
- [4] Kenneth D Harris and Thomas D Mrsic-Flogel. Cortical connectivity and sensory coding. *Nature*, 503(7474):51–58, 2013.
- [5] Mike Hawrylycz, Lydia Ng, Damon Page, John Morris, Chris Lau, Sky Faber, Vance Faber, Susan Sunkin, Vilas Menon, Ed Lein, et al. Multi-scale correlation structure of gene expression in the brain. *Neural Networks*, 24(9):933–942, 2011.
- [6] Chunlin Ji, Daniel Merl, Thomas B Kepler, and Mike West. Spatial mixture modelling for unobserved point processes: Examples in immunofluorescence histology. *Bayesian analysis (Online)*, 4(2):297, 2009.
- [7] Athanasios Kottas and Bruno Sansó. Bayesian mixture modeling for spatial poisson process intensities, with applications to extreme value analysis. *Journal of Statistical Planning and Inference*, 137(10):3151–3163, 2007.
- [8] Alex Kulesza and Ben Taskar. Determinantal point processes for machine learning. *Machine Learning*, 5(2-3):123–286, 2012.
- [9] Ed S Lein, Michael J Hawrylycz, Nancy Ao, Mikael Ayres, Amy Bensinger, Amy Bernard, Andrew F Boe, Mark S Boguski, Kevin S Brockway, Emi J Byrnes, et al. Genome-wide atlas of gene expression in the adult mouse brain. *Nature*, 445(7124):168–176, 2007.
- [10] Henry Markram, Maria Toledo-Rodriguez, Yun Wang, Anirudh Gupta, Gilad Silberberg, and Caizhi Wu. Interneurons of the neocortical inhibitory system. *Nat Rev Neurosci*, 5(10):793–807, October 2004.

- [11] Alexander Smola and Shravan Narayanamurthy. An architecture for parallel topic models. *Proceedings of the VLDB Endowment*, 3(1-2):703–710, 2010.
- [12] Amit Zeisel, Ana B Muñoz-Manchado, Simone Codeluppi, Peter Lönnerberg, Gioele La Manno, Anna Juréus, Sueli Marques, Hermany Munguba, Liquun He, Christer Betsholtz, et al. Cell types in the mouse cortex and hippocampus revealed by single-cell rna-seq. *Science*, 347(6226):1138–1142, 2015.

# Appendix for Discovering Neuronal Cell Types and Their Gene Expression Profiles Using a Spatial Point Process Mixture Model

## A Morphological Basis Extraction

We aim to characterize the morphological basis for all cells with different size, orientation, expression profiles and spatial distribution. The traditional sparse coding introduces too many free parameters and is not suitable for compact morphological basis learning. We instead propose Gaussian prior convolutional sparse coding (GPCSC). The intuition for using convolution is due to the frequent replication of cells of similar shapes and the translation invariance property. Traditional sparse coding would learn both the shape of the cell and the location of the cell. But the convolutional sparse coding would only learn the shape here. We characterize cell spatial distribution via decoding the sparse activation map.

To formulate the problem formally: let  $I$  be the image observed, then the convolutional sparse coding model generates observed image  $I$  using filters (resembling cell shapes)  $F$  superposed at locations indicated by the activation map  $M$  (whose sparsity pattern indicates cell spatial distribution and activation amplitude indicates gene expression profiles.)

Our goals of segmenting cells, extracting cell basis, and estimating gene profiles and cell locations are reduced to this optimization learning problem:

$$\min_{F_m, M_m^n} \left\| \sum_n I^n - \sum_{m=1}^k F_m \star M_m^n \right\|_F^2 + \sum_n \sum_m \lambda \|M_m^n\|_0, \text{ s.t. } F_m(x, y) \geq 0, \|F_m\|_F^2 = 1, M_m^n(x, y) \geq 0. \quad (5)$$

where  $I^n$  is the  $n^{\text{th}}$  image associated with the gene we are interested in with  $D_x \times D_y$  pixels, i.e.,  $I^n \in \mathbb{R}^{D \times D}$ .

We call the  $F_m \in \mathbb{R}^{d \times d}$  filter, where  $d$  is set to capture the local cell morphological information. The spatial coefficient for image  $I^n$  is denoted as  $H_m^n \in \mathbb{R}^{(D-d+1) \times (D-d+1)}$  which represents the position of the filter  $F_m$  being active on image  $I^n$ . More precisely, if  $H_m^n(x, y) = 1$ , then  $F_m$  is active at  $I^n(x : x + d - 1, y : y + d - 1)$ .

### A.1 Gaussian Prior Convolutional Sparse Coding

The popular alternating approach between matching pursuit to learn activation map  $M$  and k-SVD to learn  $F$  is general applicable to any object detection problem in image processing. However, this approach causes inexact cell number estimation as filters with multi-modality (i.e., multiple cells) are learnt. We resolve this issue by proposing an Gaussian probability density function prior on the filters to guarantee single cell detection and achieve accurate cell number estimation. The support of  $M$  is also limited to the local maxima indicating cell centers. Note that our cell are not donut shaped, and it is reasonable to assume the darkest point being the cell center.

Therefore, we optimize over the objective  $\min \|\sum_n I^n - \sum_m F_m \star M_m^n\|_2^2 + \sum_n \sum_m \lambda \|M_m^n\|_0$  such that  $F_m$  are 2-D Gaussian densities with priori set top 2 principal radius and orientation. Alternating Minimization is used to solving the optimization problem. If we define the residual as  $\sum_n I^n - \sum_n \sum_m \hat{F}_m \star \hat{M}_m^n$ , the gradient of the objective reduced to an iterative approach of updating filters, compute residual, optimizing activation map based on residual, compute residual and updating filters again. It is easy to see that both  $\frac{\partial L}{\partial F_m}(i, j)$  and  $\frac{\partial L}{\partial H_m}(i, j)$  are convolution of the residual and the other variable rotated by angle  $\pi$ .

### A.2 Image Registration/Alignment

A structure represents a neuronanatomical region of interest. Structures are grouped into ontologies and organized in a hierarchy or structure graph. We are interested in the somatosensory cortex area. So we use the affine transform from Allen Brain Institute [1, 9] to align all the in-situ hybridization images with the Atlas brain to extract the correct region.